arXiv:1603.02745v1 [stat.ME] 9 Mar 2016

Non-parametric latent modeling and network clustering

François Bavaud

Abstract The paper exposes a non-parametric approach to latent and co-latent modeling of bivariate data, based upon alternating minimization of the Kullback-Leibler divergence (EM algorithm) for complete log-linear models. For categorical data, the iterative algorithm generates a soft clustering of both rows and columns of the contingency table. Well-known results are systematically revisited, and some variants are presumably original. In particular, the consideration of square contingency tables induces a clustering algorithm for weighted networks, differing from spectral clustering or modularity maximization techniques. Also, we present a co-clustering algorithm applicable to HMM models of general kind, distinct from the Baum-Welch algorithm. Three case studies illustrate the theory.

1 Introduction: parametric and non-parametric mixtures

Two variables can be dependent, yet conditionally independent given a third one, that is $X \perp Y|G$ but $X \not\perp Y$: in bivariate *latent models* of dependence M, joint bivariate probabilities $P(x,y)$ express as

$$P(x,y) = \sum_{g=1}^{m} p(x,y,g) = \sum_{g=1}^{m} p(g)p(x|g)p(y|g) \tag{1}$$

where x, y, g denote the values of X, Y, G, and $p(x,y,g)$ their joint probability.

Bivariate data, such as summarized by normalized contingency tables $F(x,y) = \frac{n(x,y)}{n(\bullet,\bullet)}$, where $n(x,y)$ counts the number of individuals in $x \in X$ and $y \in Y$, can be approached by latent modeling, consisting in inferring a suitable model $P(x,y) \in M$ of

François Bavaud
University of Lausanne, e-mail: fbavaud@unil.ch

the form (1), typically closest to the observed frequencies $F(x,y)$ in the maximum-likelihood sense, or in the least squares sense. Mixture (1) also defines *memberships* $p(g|x) = p(x|g)p(g)/p(x)$ and $p(g|y)$; hence latent modeling also performs *model-based clustering*, assigning observations x and y among groups $g = 1,\ldots,m$.

Latent modeling and clustering count among the most active data-analytic research trends of the last decades. The literature is simply too enormous to cite even a few valuable contributions, often (re-)discovered independently among workers in various application fields. Most approaches are parametric, typically defining $p(x|g)$ and $p(y|g)$ as exponential distributions of some kind, such as the multivariate normal (continuous case) or the multinomial (discrete case) (see e.g. Govaert and Nadif 2013 and references therein). Parametric modelling allows further hyperparametric Bayesian processing, as in latent Dirichlet allocation (Blei et al. 2003).

By contrast, we focus on non-parametric models specified by the whole family of log-linear *complete models* $\mathscr{M}$ corresponding to $X \perp Y|G$, namely (see e.g. Christensen 2006)

$$\mathscr{M} = \{p \mid \ln p(x,y,g) = a(x,g) + b(y,g) + c\}$$

Equivalently,

$$\mathscr{M} = \{p \mid p(x,y,g) = \frac{p(x,\bullet,g)\, p(\bullet,y,g)}{p(\bullet,\bullet,g)}\}$$

where "$\bullet$" denotes the summation over the replaced argument. The corresponding class of bivariate models M of the form (1) simply reads $M = \{P \mid P(x,y) = \sum_g p(x,y,g) \equiv p(x,y,\bullet)$, for some $p \in \mathscr{M}\}$.

Observations consist of the joint empirical distribution $F(x,y)$, normalized to $F(\bullet,\bullet) = 1$. In latent modeling, one can think of the observer as a color-blind agent perceiving only the margin $f(x,y,\bullet)$ of the complete distribution $f(x,y,g)$, but not the color (or group) g itself (see Fig. 1). Initially, any member f of the set

$$\mathscr{D} = \{f \mid f(x,y,\bullet) = F(x,y)\}$$

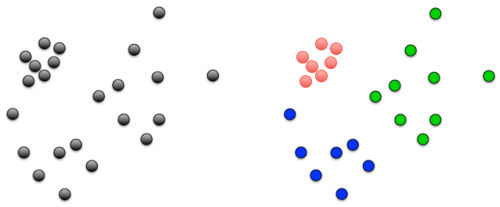

Fig. 1 Left: observed data, where (x,y) are the object coordinates. Right: complete data (x,y,g), where the group g is labeled by a color. In psychological terms, (x,y) is the *stimulus*, and (x,y,g) the *percept*, emphasizing the EM-algorithm as a possible model for cognition.

seems equally compatible with the observations F, and the role of a clustering algorithm precisely consists in selecting a few good candidates $f \in \mathscr{D}$, or even a unique one, bringing color to the observer.

This paper exposes a non-parametric approach to latent and co-latent modeling of bivariate data, based upon alternating minimization of the Kullback-Leibler divergence (EM algorithm) for complete log-linear models (section 2). For categorical data, the iterative algorithm generates a soft clustering of both rows and columns of the contingency table. Well-known results are systematically revisited, and some variants are presumably original. In particular, the consideration of square contingency tables induces a clustering algorithm for weighted networks, differing from spectral clustering or modularity maximization techniques (section 3). Also, we present a co-clustering algorithm applicable to HMM models of general kind, distinct from the Baum-Welch algorithm. Three case studies illustrate the theory: latent (co-)betrayed clustering of a term-document matrix (section 2.3), latent clustering of spatial flows (section 3.2), and latent co-clustering of bigrams in French (section 3.4).

2 EM latent clustering: a concise derivation from first principles

The *alternating minimisation* procedure (Csiszár and Tusnády 1984) provides an arguably elegant derivation of the EM algorithm; see also e.g. Cover and Thomas (1991) or Bavaud (2009). The maximum likelihood model $\hat{P} \in M$ of the form (1) minimizes the Kullback-Leibler divergence $K()$

$$\hat{P} = \arg\min_{P \in M} K(F\|P) \qquad\qquad K(F\|P) = \sum_{x,y} F(x,y) \ln \frac{F(x,y)}{P(x,y)}$$

where $F(x,y)$ denotes the empirical bivariate distribution. On the other hand, the complete Kullback-Leibler divergence $K(f\|p) = \sum_{x,y,g} f(x,y,g)\ \ln \frac{f(x,y,g)}{p(x,y,g)}$, where $f(x,y,g)$ is the empirical "complete" distribution (see fig. 1), enjoys the following properties (see e.g. Bavaud (2009) for the proofs, standard in Information Theory):

$$\hat{p}(x,y,g) := \arg\min_{p \in \mathscr{M}} K(f||p) = \frac{f(x,\bullet,g)\, f(\bullet,y,g)}{f(\bullet,\bullet,g)} \qquad \textbf{M-step} \qquad (2)$$

$$\tilde{f}(x,y,g) := \arg\min_{f \in \mathscr{D}} K(f||p) = \frac{p(x,y,g)}{p(x,y,\bullet)}\, F(x,y) \qquad \textbf{E-step} \qquad (3)$$

Furthermore, $\min_{f \in \mathscr{D}} K(f||p) = K(F||P)$, and thus

$$\min_{P \in M} K(F||P) = \min_{p \in \mathscr{M}} \min_{f \in \mathscr{D}} K(f||p)$$

Hence, starting from some complete model $p^{(0)} \in \mathscr{M}$, the EM-sequence $f^{(t+1)} := \tilde{f}[p^{(t)}]$ defined in (3) and $p^{(t+1)} := \hat{p}[f^{(t+1)}]$ defined in (2) converges towards a *local minimum* of $K(f||p)$. Observe the margins to coincide after a single EM-cycle in the sense $p^{(t)}(x,\bullet,\bullet) = F(x,\bullet)$ and $p^{(t)}(\bullet,y,\bullet) = F(\bullet,y)$ for all $t \geq 1$.

For completeness sake, note that $\mathscr{D}$ and $\mathscr{M}$ are *closed* in the following sense, as they are in other instances of the EM algorithm in general. Critically and crucially:

i) $\mathscr{D}$ is convex, that is closed under additive mixtures $\lambda f_1 + (1-\lambda) f_2$; this turns out to be the case for maximum entropy problems in general.

ii) $\mathscr{M}$ is log-convex, that is closed under multiplicative mixtures $p_1^{\lambda} p_2^{(1-\lambda)} / Z(\lambda)$ where $Z(\lambda)$ is a normalization constant; this is the case for exponential models, as well as for non-parametric log-linear models in general.

2.1 Latent co-clustering

Co-clustering describes the situation where each of the observed variables is attached to a *distinct* latent variable, the latter being mutually associated. That is, $X \perp Y|(U,V)$, $X \perp V|U$ and $Y \perp U|V$ while $X \not\perp Y$, and $U \not\perp V$ in general. Equivalently, $X \to U \to V \to Y$ form a "Markov chain", in the sense of Cover and Thomas (1991). Bivariate joint probabilities express as

$$P(x,y) = \sum_{u=1}^{m_1} \sum_{v=1}^{m_2} p(x,y,u,v) = \sum_{u,v} p(u,v) p(x|u) p(y|v) \tag{4}$$

Complete models $\mathscr{M}$, restricted models M and complete empirical distributions $\mathscr{D}$ are

$$\mathscr{M} = \{p \,|\, p(x,y,u,v) = \frac{p(x\bullet u\bullet)\, p(\bullet y\bullet v)\, p(\bullet\bullet uv)}{p(\bullet\bullet u\bullet)\, p(\bullet\bullet\bullet v)}\} \tag{5}$$

$$M = \{P \,|\, P(x,y) = p(x,y,\bullet,\bullet) \text{ with } p \in \mathscr{M}\} \tag{6}$$

$$\mathscr{D} = \{f \,|\, f(x,y,\bullet,\bullet) = F(x,y)\} \tag{7}$$

where $F(x,y)$ denotes the observed empirical distribution. The steps of the former section apply again, yielding the EM algorithm

$$\hat{p}(x,y,u,v) := \arg\min_{p \in \mathscr{M}} K(f||p) = \frac{f(x\bullet u\bullet)\, f(\bullet y\bullet v)\, f(\bullet\bullet uv)}{f(\bullet\bullet u\bullet)\, f(\bullet\bullet\bullet v)} \qquad \textbf{M-step} \tag{8}$$

$$\tilde{f}(x,y,u,v) := \arg\min_{f \in \mathscr{D}} K(f||p) = \frac{p(x,y,u,v)}{p(x,y,\bullet,\bullet)}\, F(x,y) \qquad \textbf{E-step} \tag{9}$$

where $K(f\|p) = \sum_{x,y,u,v} f(x,y,u,v) \ln \frac{f(x,y,u.v)}{p(x,y,u,v)}$ measures the divergence of the complete observations from the complete model.

2.2 Matrix and tensor algebra for contingency tables

The material of sections (2) and (2.1) holds *irrespectively of the continuous or discrete nature of X and Y*: in the continous case, integrals simply replace sums. In the discrete setting, addressed here, categories are numbered as $i = 1,\ldots,n$ for X, as $k = 1,\ldots,p$ for Y and as $g = 1,\ldots,m$ for G. Data consist of the relative $n \times p$ contingency table F_{ik} normalized to $F_{\bullet\bullet} = 1$.

2.2.1 Latent co-clustering

Co-clustering models and complete models express as

$$P_{ik} = \sum_{u=1}^{m_1}\sum_{v=1}^{m_2} c_{uv}\, a_i^u\, b_k^v \qquad\qquad p_{ikuv} = c_{uv}\, a_i^u\, b_k^v \tag{10}$$

- where $c_{uv} = P(U = u, V = v) = p(\bullet\bullet uv)$, obeying $c_{\bullet\bullet} = 1$, is the *joint latent distribution* of row, respectively column groups u and v
- $a_i^u = p(i\bullet u\bullet)/p(\bullet\bullet u\bullet)$ (with $a_\bullet^u = 1$) is the *row distribution conditionally to the row group* $U = u$, also referred to as *emission probability* (section 3)
- $b_k^v = p(\bullet k\bullet v)/p(\bullet\bullet\bullet v)$ (with $b_\bullet^v = 1$) is the column distribution or emission probability conditionally to the column group $V = v$.

Hence, a complete model p is entirely determined by the triple (C,A,B), where $C = (c_{uv})$ is $m_1 \times m_2$ and normalized to unity, $A = (a_i^u)$ is $n \times m_1$ and $B = (b_k^v)$ is $p \times m_2$, both row-standardized.

It is straightforward to show that the successive application of the E-step (9) and the M-step (8) to $p \equiv (C,A,B)$ yields the new complete model $\ddot{p} \equiv (\ddot{C},\ddot{A},\ddot{B})$ with

$$\ddot{c}_{uv} = c_{uv} \sum_{jl} \frac{F_{jl}}{P_{jl}}\, a_j^u\, b_l^v \tag{11}$$

$$\ddot{a}_i^u = a_i^u \frac{\sum_{lv'} c_{uv'} \frac{F_{il}}{P_{il}}\, b_l^{v'}}{\sum_{jlv'} c_{uv'} \frac{F_{jl}}{P_{jl}}\, a_j^u\, b_l^{v'}} \tag{12}$$

$$\ddot{b}_k^v = b_k^v \frac{\sum_{ju'} c_{u'v} \frac{F_{jk}}{P_{jk}}\, a_j^{u'}}{\sum_{jlu'} c_{u'v} \frac{F_{jl}}{P_{jl}}\, a_j^{u'}\, b_l^v} \tag{13}$$

Also, after a single EM cycle, margins are respected, that is $\ddot{P}_{i\bullet} = F_{i\bullet}$ and $\ddot{P}_{\bullet k} = F_{\bullet k}$.

In hard clustering, rows i are attached to a single group denoted $u[i]$, that is $a_i^u = 0$ unless $u = u[i]$; similarly, $b_k^v = 0$ unless $v = v[k]$. Restricting P in (10) to hard clustering yields *block clustering*, for which $K(F||P) = I(X:Y) - I(U:V)$, where

$I()$ is the mutual information (e.g. Kullback (1959); Bavaud (2000); Dhillon et al. (2003)).

The set M of models P of the form (10) is convex, with extreme points consisting of hard clusterings. $K(F\|P)$ being convex in P, its minimum is attained for convex mixtures of hard clusterings, that is for *soft clusterings*.

2.2.2 Latent clustering

Setting $m_1 = m_2 = m$ and C diagonal with $c_{gh} = \rho_g \, \delta_{gh}$ yields the latent model

$$P_{ik} = \sum_{g=1}^{m} \rho_g \, a_i^g \, b_k^g \qquad\qquad p_{ikg} = \rho_g \, a_i^g \, b_k^g \tag{14}$$

together with the corresponding EM-iteration $p \equiv (\rho, A, B) \to \ddot{p} \equiv (\ddot{\rho}, \ddot{A}, \ddot{B})$, namely

$$\ddot{\rho}_g = \rho_g \, \kappa_g \qquad\qquad \ddot{a}_i^g = a_i^g \frac{\sum_l b_l^g \frac{F_{il}}{P_{il}}}{\kappa_g} \qquad\qquad \ddot{b}_k^g = b_k^g \frac{\sum_j a_j^g \frac{F_{jk}}{P_{jk}}}{\kappa_g} \tag{15}$$

where $\kappa_g = \sum_{jl} a_j^g b_l^g \frac{F_{jl}}{P_{jl}}$. Similar, if not equivalent updating rules have been proposed in information retrieval and natural language processing (Saul and Pereira 1997; Hofmann 1999), as well as in the *non-negative matrix factorization* framework (Lee and Seung 2001; Finesso and Spreij 2006).

By construction, families of latent models (14) M_m with m groups are nested in the sense $M_m \subseteq M_{m+1}$.

The case $m = 1$ amounts to *independence models* $P_{ik} = a_i b_k$, for which the fixed point $\ddot{a}_i = F_{i\bullet}$ and $\ddot{b}_k = F_{\bullet k}$ is, as expected, reached after a single iteration, irrespectively of the initial values of a and b.

By contrast, $m \geq \mathrm{rank}(F)$ generates *saturated models*, exactly reproducing the observed contingency table. For instance, assume that $m = p = \mathrm{rank}(F) \leq n$; then taking $a_i^g = F_{ig}/F_{\bullet g}$, $b_k^g = \delta_{kg}$ and $\rho_g = F_{\bullet g}$ (which already constitutes a fixed point of (15)) evidently satisfies $P_{ik} = F_{ik}$.

2.3 Case study I: Reuters 21578 term-document matrix

The $n \times p = 20 \times 1266$ document-term normalized matrix F, constituting the *Reuters 21578* dataset, is accessible through the R package **tm** (Feinerer et al. 2008). The co-clustering algorithm (11) (12) (13) is started by randomly assigning uniformly each document to a single row group $u = 1, \ldots, m_1$, and by uniformly assigning each term to a single column group $v = 1, \ldots, m_2$. The procedure turns out to converge

after about 1000 iterations (figure 2), yielding a locally minimal value $K_{m_1 m_2}$ of the Kullback-Leibler divergence. By construction, $K_{m_1 m_2}$ decreases with m_1 and m_2. Latent clustering (15) with m groups is performed analogously, yielding a locally minimal value K_m.

Experiments with three or four groups yield the typical results $K_3 = 1.071180 > K_{33} = 1.058654 > K_{43} = 1.038837 > K_{34} = 1.036647 > K_4 = 0.877754 > K_{44} = 0.873071$. The above ordering is expected, although inversions are frequently observed, under differing random initial configurations. Model selection procedures, not addressed here, should naturally consider in addition the degrees of freedom, larger for co-clustering models. The latter do not appear as particularly rewarding here (at least for the experiments performed, and in contrast to the results associated to case study III of section 3.4): indeed, joint latent distributions C turn out to be "maximally sparse", meaning that row groups u and column groups v are essentially the same. Finally, each of the 20 documents of the Reuters 2157 dataset happens to belong to a single row group (hard clusters), while only a minority of the 1266 terms (say about 20%) belong to two or more column groups (soft clusters).

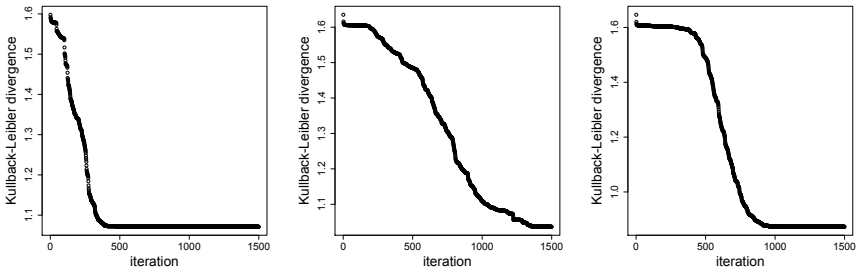

Fig. 2 Case study I: convergence of the latent and co-latent iterating procedure. Left: latent model with $m = 3$. Middle: co-latent model with $(m_1, m_2) = (3,4)$. Right: co-latent model with $(m_1, m_2) = (4,4)$.

3 Network clustering

When the observed categories x and y belong to the same set indexed by $i, j = 1, \ldots, n$, the relative square contingency table F_{ij} defines a *directed weighted network* on n vertices: F_{ij} is the weight of edge (ij), $F_{i\bullet}$ is the outweight of vertex i (relative outdegree) and $F_{\bullet i}$ its inweight i (relative indegree), all normalized to unity. Frequently, F_{ij} counts the relative number of units initially at vertex i, and at vertex j after some fixed time. Examples abound in spatial migration, spatial commuting, social mobility, opinion shifts, confusion matrices, textual dynamics, etc.

A further restriction, natural in many applications of latent network modeling, consists in identifying the row and column emission probabilities, that is in requiring

$b_i^g = a_i^g$. This condition generates four families of nested latent network models of increasing flexibility, namely

$$P_{ij} = \sum_{g=1}^{m} \rho_g \, a_i^g \, a_j^g \qquad \text{latent (symmetric) network model} \quad (16)$$

$$P_{ij} = \sum_{u,v=1}^{m} c_{uv} \, a_i^u \, a_j^v \quad \text{with } c_{uv} = c_{vu} \qquad \text{co-latent symmetric network model} \quad (17)$$

$$P_{ij} = \sum_{u,v=1}^{m} c_{uv} \, a_i^u \, a_j^v \quad \text{with } c_{u\bullet} = c_{\bullet u} \qquad \text{co-latent MH network model} \quad (18)$$

$$P_{ij} = \sum_{u,v=1}^{m} c_{uv} \, a_i^u \, a_j^v \qquad \text{co-latent general network model} \quad (19)$$

Models (16) and (17) $P = P'$, making latent and co-latent *symmetric* clustering suitable for *unoriented weighted networks* with $F_{ij} = F_{ji}$. By contrast, unrestricted co-latent models (19) describes general *oriented weighted networks*. Symmetric matrices $F = (F_{ij})$ appear naturally in reversible random walks on networks, or in spatial modeling where they measure the spatial interaction between regions (spatial weights), and constitute a weighted version of the adjacency matrix, referred to as an *exchange matrix* by the author (Bavaud 2014 and references therein; see also Berger and Snell 1957).

Latent models (16) are positive semi-definite or *diffusive*, that is endowed with *non-negative eigenvalues*, characteristic of a continuous propagation process from a one place to its neighbours. In particular, the diagonal part of P in (16) cannot be too small. In contrast, co-latent symmetrical network models (17) are flexible enough to describe phenomena such as bipartition or periodic alternation, implying negative eigenvalues.

The condition (18) of *marginal homogeneity* (MH) on the joint latent distribution C is inherited by the restricted models, in the sense $P_{i\bullet} = P_{\bullet i}$. They constitute appropriate models for the bigram distributions of single categorical sequences (of length N, constituted of n types), for which $F_{i\bullet} = F_{\bullet i} + O(N^{-1})$; see the case study III of Section 3.4. Formulation (18) describes m hidden states related by a Markov transition matrix $p(v|u) = c_{uv}/c_{u\bullet}$, as well as n observed states related to the hidden ones by the *emission probabilities* $a_i^u = p(i|u)$. Noticeably enough, (18) precisely encompasses the ingredients of the *hidden Markov models* (HMM) (see e.g. Rabiner 1989).

3.1 Network latent clustering

Approximating F by P in (16) amounts in performing a *soft network clustering*: the *membership* of vertex i in group g (of weight ρ_g) is

$$z_{ig} = p(i|g) = \frac{p(i)p(g|i)}{p(g)} = \frac{f_i a_i^g}{\rho_g} \qquad \text{with} \quad f_i = F_{i\bullet} = F_{\bullet i} \quad \text{and} \quad \rho_g = \sum_{i=1}^{n} f_i z_{ig} \ .$$

EM-updating rules for memberships (instead of emission probabilities, for a change)

$$P_{ij} = f_i f_j \sum_{g=1}^{m} \frac{z_{ig} z_{jg}}{\rho_g} \qquad \ddot{z}_{ig} = z_{ig} \sum_j \frac{F_{ij}}{P_{ij}} \frac{f_j z_{jg}}{\rho_g} \qquad \ddot{\rho}_g = \sum_i f_i \ddot{z}_{ig} \qquad (20)$$

define a soft clustering iterative algorithm for unoriented weighted networks, presumably original.

3.2 Case study II: inter-cantonal Swiss migrations

Consider the $n \times n$ matrix $N = (N_{ij})$ of inter-cantonal migratory flows in Switzerland, counting the number of people inhabiting canton i in 1980 and canton j in 1985, $i, j = 1, \ldots, n = 26$, for a total of $\text{sum}(N) = 6'039'313$ inhabitants, 93% of which lie on the diagonal (stayers). The symmetric, normalized matrix $F = \frac{1}{2}(N + N')/N_{\bullet\bullet}$ is diffusive, largely dominated by its diagonal. As a consequence, direct application of algorithm (20) from an initial random cantons-to-groups assignation produces somewhat erratic results: a matrix $F = (F_{ij})$ too close to the identity matrix $I = (\delta_{ij})$ cannot by reasonably approximated by the latent model (16), unless $m = n$, where each canton belongs to its own group.

Here, the difficulty lies in the shortness of the observation period (5 years, smaller than the average moving time), making the off-diagonal contribution $1 - \text{trace}(F)$ too small. Multiplying the observation period by a factor $\lambda > 1$ generates, up to $O(\lambda^2)$, a modified relative flow $\tilde{F}_{ij} = \lambda F_{ij} + (1-\lambda)\delta_{ij} f_i$, where $f_i = F_{i\bullet} = F_{\bullet i}$ is the weight of canton i. The modified $\tilde{F}$ is normalized, symmetric, possesses unchanged vertex weights $\tilde{F}_{i\bullet} = f_i$, and its off-diagonal contribution is multiplied by λ. Of course, λ cannot be too large, in order to insure the non-negativity of $\tilde{F}$ ($\lambda \leq 6.9$ here) as well as its semi-positive definiteness ($\lambda \leq 6.4$ here).

Typical realizations of (20), with $\lambda = 5$, are depicted in figures (3) and (4): as expected, spatially close regions tend to be regrouped.

3.3 Network general co-clustering

Latent co-clustering (19) applies to contingency tables F of general kind, possibly asymmetric or marginally inhomogeneous, and possibly exhibiting diffusivity, alternation, or a mixture of them. Implementing the common emission probabilities

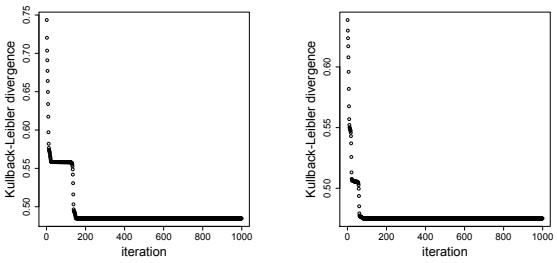

Fig. 3 Decrease of the Kullback-Leibler divergence for the two realizations of figure 4, respectively. Horizontal plateaux correspond to metastable minima in the learning of the latent structure, followed by the rapid discovery of a better fit.

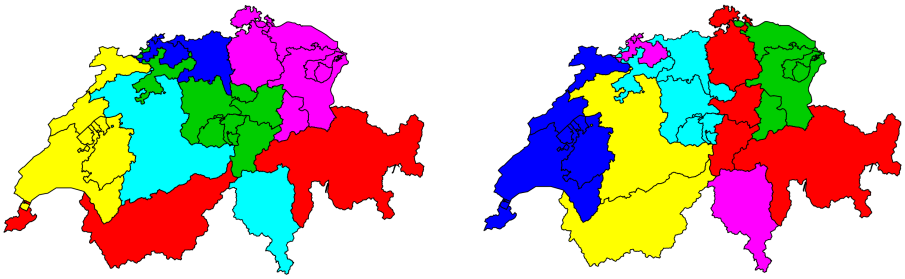

Fig. 4 Case study II: two realizations of the network latent clustering algorithm (20), applied to the modified flow matrix $\tilde{F}$, with random initial assignment to $m = 6$ groups, and final hard assignment of canton i to group $\arg\max_g z_{ig}$.

constraint in the M-step (8) yields together with (19) the updating rule

$$\ddot{c}_{uv} = c_{uv} \sum_{ij} \frac{F_{ij}}{P_{ij}}\, a_i^u\, a_j^v \qquad\qquad \ddot{a}_i^u = a_i^u\, \frac{\sum_{j'v'} (c_{uv'}\, \frac{F_{ij'}}{P_{ij'}} + c_{v'u}\, \frac{F_{j'i}}{P_{j'i}})\, a_{j'}^{v'}}{\sum_{i'j'v'} (c_{uv'}\, \frac{F_{i'j'}}{P_{i'j'}} + c_{v'u}\, \frac{F_{j'i'}}{P_{j'i'}})\, a_{i'}^u\, a_{j'}^{v'}} \tag{21}$$

Let us recall that the classical Baum-Welch algorithm handles the HMM modeling of a single (large) sequence of tokens, with (almost) marginally homogeneous bigram counts. By contrast, the presumably original iterative algorithm (21) also seems to be able to handle marginally inhomogeneous network data, such that aggregates of smaller sequences. Further experimentations are needed at this stage to gauge the generality of model (19), and the efficiency of algorithm (21).

For symmetric data $F = F'$, the symmetric model (17) can be tackled by (21) above, with the simplifying circumstance that the additive symmetrizing occurring in the numerator and denominator of $\ddot{a}_i^u$ is not needed anymore, provided that the initial joint probability c_{uv} is symmetrical, which automatically insures the symmetry of further iterates $\ddot{c}_{uv}$.

3.4 Case study III: modeling bigrams

group	1	2	3	4
_			61	
a				25
b	3			
c	10	1	1	
d	12	2		
e	1			45
f	4			
g	3	1		
h	2			1
i	1	3	13	5
j	2			
k				
l	17			4
m	8		2	
n	17	9		
o				13
p	8		1	
q		5		
r	10	6	4	3
s	6	28		
t	7	26		
u		8	10	3
v	7			
w				
x		2		
y	1			
z		1		

group	1	2	3	4
_			100	
a				100
b	89		11	
c	87	5	8	
d	90	10		
e	2		1	97
f	87			13
g	71	26		3
h	65			35
i	4	9	61	26
j	100			
k		100		
l	77	2		21
m	76		24	
n		53	47	
o				100
p	89		9	2
q		88		12
r	40	21	22	17
s	20	80		
t	26	74		
u		27	57	16
v	100			
w	100			
x		100		
y	60	40		
z		100		

	1	2	3	4
1	0	0	1	21
2	0	3	12	2
3	17	7	0	7
4	5	8	17	0

	1	2	3	4
1	0	0	7	93
2	0	15	71	14
3	54	24	0	22
4	17	27	57	0

1	22
2	18
3	31
4	30

Table 1 Case study III: emission probabilities A (left), memberships Z (middle), joint latent distribution C (right, top), latent probability transition matrix W (right, middle) and its corresponding stationary distribution π (right, bottom). All values are multiplied by 100 and rounded to the nearest integer.

We consider the first chapters of the French novel "La Bête humaine" by Emile Zola (1890). After suppressing all punctuation, accents and separators with exception of the blank space, and converting upper-case letters to lower-case, we are left with a sequence of $N = 725'000$ tokens on $n = 27$ types (the alphabet + the space), containing 724'999 pairs of successive tokens or *bigrams*. The resulting $n \times n$ normalized contingency table $F = (F_{ij})$ is far from symmetric (for instance, the bigram **qu** occurs 6'707 times, while **uq** occurs only 23 times), but almost *marginally homogenous*, that is $F_{i\bullet} \cong F_{\bullet i} + 0(N^{-1})$ (and exactly marginally homogenous if one starts and finishes the textual sequence with the same type, such as a blank space).

Symmetrizing F as $F^s = (F + F')/2$ does not makes it diffusive, and hence unsuitable by latent modelling (16), because of the importance of large negative eigenvalues in F^s, betraying *alternation*, typical in linguistic data - think in particular of the vowels-consonants alternation (e.g. Goldsmith and Xanthos 2009). This being said, symmetric co-clustering of F^s (17) remains a possible option.

Table 1 results from the general co-clustering algorithm (21) applied on the original, asymmetric bigram counts F itself. Group 4 mainly emits the vowels, group 3 the blank, group 2 the `s` and `t`, and group 1 other consonants. Alternation is betrayed by the null diagonal of the Markov transition matrix W - with the exception of group 2.

The property of marginal homogeneity $F_{i\bullet} = F_{\bullet i}$ permits in addition to obtain the *memberships Z from the emissions A*, by first determining the solution ρ of $\sum_g \rho_g a_i^g = f_i$, where $f_i = F_{i\bullet} = F_{\bullet i}$ is the relative frequency of letter i, and then by defining $z_{ig} = \rho_g a_i^g / f_i$.

References

1. Bavaud, F.: An Information Theoretical approach to Factor Analysis. In Proceedings of the 5th International Conference on the Statistical Analysis of Textual Data (JADT 2000), pp. 263–270 (2000)
2. Bavaud, F.: Information theory, relative entropy and statistics. In Sommaruga, G. (Ed.) Formal Theories of Information, LNCS 5363 pp. 54–78, Springer (2009)
3. Bavaud, F.: Spatial weights: constructing weight-compatible exchange matrices from proximity matrices. In Duckham, M. et al. (Eds.) GIScience 2014, LNCS 8728, pp. 81–96, Springer (2014)
4. Berger, J., Snell, J. L.: On the concept of equal exchange. Systems Research and Behavioral Science **2**, pp. 111–118 (1957)
5. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent Dirichlet allocation. The Journal of machine Learning research **3**, pp. 993–1022 (2003)
6. Christensen, R: Log-linear models and logistic regression. Springer, (2006)
7. Cover, T. M., Thomas, J. A.: Elements of Information Theory. Wiley (1991)
8. Csiszár, I., Tusnády, G.: Information Geometry and Aternating Minimization Procedures. In: Dedewicz, E.F. (ed.) Statistics and Decisions, Supplement Issue 1 pp. 205–237 (1984)
9. Dhillon, I.S., Mallela, S., Modha, D.S.: Information-theoretic co-clustering. In Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 89–98 (2003)
10. Feinerer, I., Hornik, K., Meyer, D.: Text mining infrastructure in R. Journal of Statistical Software, 25(5), pp. 1–54 (2008)
11. Finesso, L., Spreij, P. : Nonnegative matrix factorization and I-divergence alternating minimization. Linear Algebra and its Applications 416, pp. 270–287 (2006)
12. Goldsmith, J., Xanthos, A.: Learning Phonological Categories. Language 85, pp. 4–38 (2009)
13. Govaert, G., Nadif, M. Co-Clustering. Wiley (2013)
14. Hofmann, T.: Probabilistic latent semantic indexing. In Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval, pp. 50–57 (1999)
15. Kullback, S.: Information theory and statistics. Wiley (1959)
16. Lee, D. D., Seung, H. S.: Algorithms for non-negative matrix factorization. In Advances in neural information processing systems, pp. 556–562 (2001)
17. Rabiner, L. R.: A tutorial on hidden Markov models and selected applications in speech recognition. Proceedings of the IEEE, **77**(2), pp. 257–286 (1989)
18. Saul, L., Pereira, F.: Aggregate and mixed-order Markov models for statistical language processing. In Proceedings of the 2nd International Conference on Empirical Methods in Natural Language Processing (1997)